A Comparison of Multilevel Preconditioners for Solving Multimaterial Equilibrium Radiation Diffusion Problems on Locally Refined Grids

Bobby Philip, T-7; Michael Pernice, Idaho National Laboratory; Wayne Joubert, Xiylon Software; Bryan Lally, CCS-2

The purpose of this article is to compare geometric multilevel preconditioners tailored for adaptive mesh refinement (AMR) grids (FAC, AFAC, AFACx) with algebraic multigrid (AMG) methods on structured AMR (SAMR) grids.

In order to obtain locally refined grids that are representative of what is encountered in practice, the preconditioners are evaluated within the context of a fully dynamic adaptive simulation [1] of the propagation of radiation or thermal energy using a diffusion approximation.

Each time-advanced solution for the radiation diffusion problem is found using the Jacobian-free Newton-Krylov method in PETSc [3]. Each time step thus requires a sequence of linear system solves. Preconditioning these linear systems requires a diffusion solve [1] for which we compare FAC, AFAC, AFACx, and AMG. SAMRAI [4] is used to manage the complexity of dynamic locally refined grids.

Table 1 shows the iteration counts obtained for FAC. Comparable iteration count results were also obtained for AFACx and LAMG (see [2]).

Table 1: *Summary of FAC iterations.*

J	1	2	3	4	5
32 x 32	--	9.7	9.8	10.3	11.6
64 x 64	--	15.1	14.8	14.1	--
128 x 128	7.3	9.0	10.4	--	--
256 x 256	7.5	10.0		--	--
512 x 512	8.3	--	--	--	--

The average number of linear iterations per time step is shown. The first column gives the size of the base grid; J *is the number of refinement levels. Performance at a fixed finest resolution is obtained by reading diagonally from lower left to upper right. Grid configurations that were not run are denoted by --.*

Figure 1 plots the wall-clock time spent in the FAC preconditioner for different base grid configurations. Similar results were obtained for AFACx (see [2]). The smaller problem sizes due to AMR result in fairly flat graphs as the number of processors is increased beyond 8. Communication costs do not appear to dominate the computations as the number of processors is increased.

Figure 2 compares the parallel performance of FAC and AFACx for a fixed AMR configuration. Similar plots are obtained for other configurations also (see [2]). AFACx was not optimized or load balanced in a manner that would exploit the asynchronous nature of the algorithm, explaining the similar performance obtained. A possible optimization in the future would be to redistribute the load within an asynchronous preconditioning step.

Figure 3 compares the parallel performance of AMG, FAC, and AFACx on two cases for which we have complete data. In general, AMG requires approximately twice the execution time as FAC, and AFACx, for the model problems tested. We note that the setup phase for AMG, which requires remapping and ordering the grid hierarchy introduces an overhead of roughly 10%. None of the approaches is especially scalable in these evaluations; this can be ascribed in part to the small problem sizes that are made possible through the use of adaptive mesh refinement.

For more information contact Bobby Philip at bphilip@lanl.gov.

[1] M. Pernice, B. Philip, *SIAM J. Sci. Comput.* **27** (5), pp. 1709–1726 (2006).
[2] M. Pernice, et al., "A Comparison of Multilevel Preconditioners for Solving Multimaterial Equilibrium Radiation Diffusion Problems on Locally Refined Grids," Los Alamos National Laboratory report LA-UR-06-7231.
[3] S. Balay, et al., "PETSc Users Manual," ANL-95/11, Argonne National Laboratory.
[4] R. Hornung, S. Kohn, *Concurrency Comput.: Pract. Exp.* **14**, pp. 347–368 (2002).

Funding Acknowledgements
NNSA's Advanced Simulation and Computing (ASC), and Weapons Supported Research (WSR) Program.

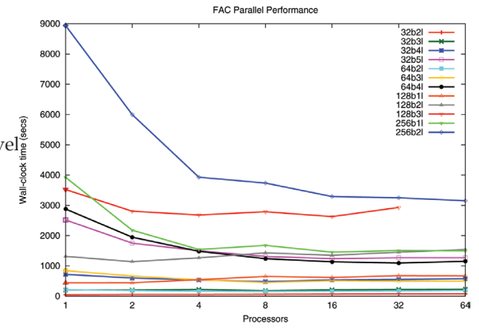

Fig. 1. *Parallel performance of FAC with varying base grid resolution and number of refinement levels.*

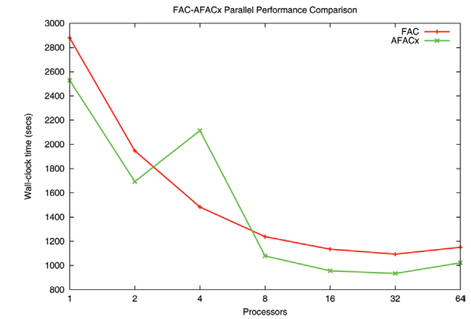

Fig. 2. *Comparison of parallel performance for FAC and AFACx for an AMR configuration with a 64 x 64 base grid and four levels of refinement.*

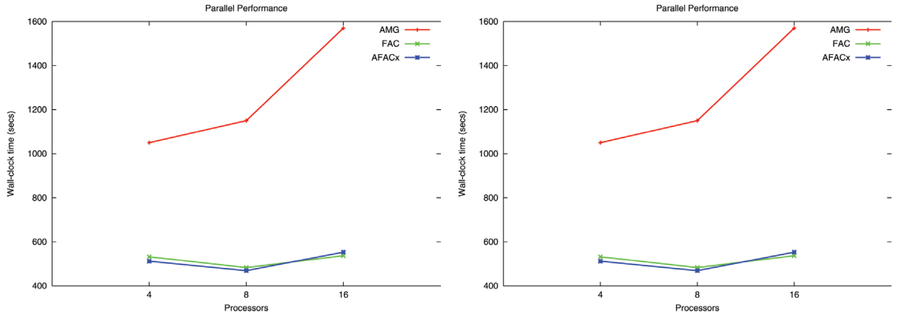

Fig. 3. *Comparison of parallel performance of the various multilevel solvers. The plot on the left is for the case with a 32 x 32 base grid and two additional refinement levels; the plot on the right has the same base grid and three refinement levels.*